arXiv:1410.5385v2 [math.DS] 10 Aug 2016

TRIANGLES IN CARTESIAN SQUARES OF QUASIRANDOM GROUPS

VITALY BERGELSON, DONALD ROBERTSON, AND PAVEL ZORIN-KRANICH

Abstract. We prove that triangular configurations are plentiful in large subsets of cartesian squares of finite quasirandom groups from classes having the quasirandom ultraproduct property, for example the class of finite simple groups. This is deduced from a strong double recurrence theorem for two commuting measure-preserving actions of a minimally almost periodic (not necessarily amenable or locally compact) group on a (not necessarily separable) probability space.

1. Introduction

By showing that any subset of $\mathrm{PSL}(2,\mathbb{F}_q)$ with density at least $2|\mathrm{PSL}(2,\mathbb{F}_q)|^{8/9}$ contains a subset of the form $\{g,x,gx\}$, Gowers [Gow08, Theorem 3.3] answered negatively a question of Babai and Sós [BS85] on the existence of a constant $c>0$ such that every finite group G has a subset of size at least $c|G|$ that is *product-free*, meaning that it contains no subset of the form $\{g,x,gx\}$. Gowers also showed [Gow08, Lemma 5.1] that quasirandom groups constitute the general setting for such a result; a finite group is *D-quasirandom* if and only if it has no non-trivial representations over $\mathbb{C}$ of dimension less than D.

The non-existence of large product-free sets in infinite, amenable groups was investigated in [BF09], where it was shown that G is minimally almost periodic (meaning that it has no non-trivial, finite-dimensional, unitary representations over $\mathbb{C}$) if and only if every subset having positive density with respect to some Følner sequence contains a subset of the form $\{g,x,gx\}$.

In both settings, the non-existence of large product-free sets is related to the absence of finite dimensional representations. It can also be related to the ergodic theory of group actions. Indeed, given $A\subset G$, one can find in A a set of the form $\{g,x,gx\}$ if and only if there is $g\in A$ such that $A\cap g^{-1}A\neq\varnothing$ and one is now faced with a question of recurrence for the action of G on itself by left multiplication. When G is finite, the Haar measure on G is a natural invariant measure, while for countable amenable groups a version of the Furstenberg correspondence principle (e.g. [Ber00, Theorem 4.17]) can be used to phrase the problem dynamically. In terms of ergodic theory, then, the question becomes one of relating the representation theory of G to positivity of correlations $\mu(B\cap(T^g)^{-1}B)$ for an action T of G on a probability space $(X,\mathscr{B},\mu)$.

With this framework in mind, the existence of more complicated configurations in subsets of quasirandom groups was considered in [BT14], and in particular the question of whether every large enough subset A of a D-quasirandom group contains a configuration of the form $\{g,x,gx,xg\}$ or, equivalently, whether $A\cap g^{-1}A\cap Ag^{-1}$ is non-empty for some $g\in A$. Dynamically this corresponds to positivity of a multiple correlation of the form

$$\mu(B\cap(T_1^g)^{-1}B\cap(T_1^gT_2^g)^{-1}B) \tag{1.1}$$

where μ is normalized counting measure and T_1,T_2 are the commuting actions of G on itself determined by left and right multiplication. It was shown, in the following strong form, that the conjugation-invariant subsets of G are the only obstruction to the positivity of such correlations.

Theorem 1.2 ([BT14, Theorem 5])**.** *Let G be a finite, D-quasirandom group with normalized Haar measure* m *and let $f_1,f_2,f_3:G\to\mathbb{R}$ be bounded in absolute value by 1. Then*

$$\int\left|\int f_1(x)\cdot f_2(xg)\cdot f_3(gx)\,\mathrm{dm}(x)-\int f_1\,\mathrm{dm}\int f_2\cdot\mathbb{E}(f_3|\mathcal{I}_G)\,\mathrm{dm}\right|\mathrm{dm}(g)\leq c(D)$$

Date: August 11, 2016.

2010 *Mathematics Subject Classification.* Primary 05D10; Secondary 28D15.

The first author gratefully acknowledges the support of the NSF under grants DMS-1162073 and DMS-1500575.

where $\mathbb{E}(f_3|\mathcal{I}_G)$ *is the orthogonal projection in* $\mathrm{L}^2(G,\mathrm{m})$ *of* f_3 *on the conjugation-invariant functions and* $c(D)$ *is a quantity depending only on* D *that goes to zero as* $D \to \infty$.

By specializing to $f_1 = f_2 = f_3 = 1_A$ where A is a subset of a D-quasirandom group G, it follows (see [BT14, Corollary 6]) that for any $\varepsilon > 0$ one has

$$\frac{|A|^3}{|G|^3} - \varepsilon \leq \frac{|A \cap gA \cap Ag|}{|G|} \leq \frac{|A|^2}{|G|^2} + \varepsilon$$

for all but at most $\varepsilon^{-1}c(D)|G|$ many $g \in G$, where $c(D) \to 0$ as $D \to \infty$. Thus if $|A| > \varepsilon^{-1}c(D)|G|$ then there $g \in A$ for which many configurations of the form $\{x, gx, xg\}$ can be found in A.

Theorem 1.2 has recently been reproved by Austin [Aus15] without the use of ultra quasirandom groups, yielding explicit bounds for $c(D)$. In particular [Aus15, Theorem 1] implies that for any D-quasirandom group G and any $A \subset G$ with $|A|^4 > 4D^{-1/8}|G|^4$ one has

$$\mathrm{m}(A \cap g^{-1}A \cap Ag^{-1}) \geq \mathrm{m}(A)^3 - \frac{4D^{-1/8}}{\mathrm{m}(A)} > 0$$

for some $g \in A$. See also [Tse14], where the bound in [Aus15] was improved and Theorem 1.2 was generalized to the setting of probability groups.

Correlations of the form (1.1) control the existence of many other types of configuration. To describe a couple, in [BMZ97] positivity of (1.1) on average for arbitrary commuting actions T_1 and T_2 of any countable, amenable group G on a probability space was proven and used to exhibit triangular configurations of the form $\{(x,y),(gx,y),(gx,gy)\}$ in any positive-density subset of $G \times G$, and in [BCRZK16] the correlation (1.1) was shown to be larger than $\mu(B)^4$ on average when T_1 and T_2 are any commuting actions of an amenable, minimally almost periodic group having the property that the $G \times G$ action $(g_1, g_2) \mapsto T_1^{g_1}T_2^{g_2}$ is ergodic. This was used to exhibit two-sided finite products sets in positive-density subsets of such groups.

In this paper we show (see Theorem 1.12 below) that there are many triangles, i.e. configurations of the form $\{(x,y),(gx,y),(gx,gy)\}$, in large enough subsets of $G \times G$ provided G is quasirandom enough. To do this we consider the behavior of the correlation

$$\int f_0(x,y) \cdot f_1(gx,y) \cdot f_2(gx,gy)\,\mathrm{d}(\mathrm{m} \times \mathrm{m})(x,y) \tag{1.3}$$

for functions $f_0, f_1, f_2 : G \times G \to \mathbb{R}$. As in [BT14] we do not work with a specific quasirandom group G directly, but instead consider the asymptotic behavior of (1.3) along quasirandom sequences of groups.

Definition 1.4. A sequence $n \mapsto G_n$ of finite, D_n-quasirandom groups is a *quasirandom* sequence of groups if $D_n \to \infty$ as $n \to \infty$.

Given a quasirandom sequence $n \mapsto G_n$ of groups, we relate the asymptotic behavior of (1.3) to a correlation of the form

$$\int f_0 \cdot T_1^g f_1 \cdot T_1^g T_2^g f_2\,\mathrm{d}\mu \tag{1.5}$$

for commuting actions T_1 and T_2 of a limiting group G formed from the G_n on a probability space $(X, \mathscr{B}, \mu)$. One can see that if f_0 and f_1 are supported on disjoint T_1-invariant sets then (1.5) is zero, so the correlation depends on the conditional expectations of f_0 and f_1 on the sub-σ-algebra of T_1-invariant sets. Similarly, the result depends on the expectations of f_1 and f_2 on the T_2-invariant sets, and on the expectations of f_0 and f_2 on the T_1T_2-invariant sets.

In order to make precise the dependence of (1.5) on the invariant sub-σ-algebras mentioned above, one studies the limiting behavior of (1.5) along some limiting scheme, a method that has been in use ever since Furstenberg's ergodic proof [Fur77] of Szemerédi's theorem. Which limiting scheme is used, and which sub-σ-algebras control the limiting behavior, depends on the properties of the acting group G.

When G is countable and amenable one can use a Følner sequence $N \mapsto \Phi_N$ to average (1.1). Austin [Aus13] has shown, using his satedness technique – see Section 3, that when $(X, \mathscr{B}, \mu)$ is a standard probability space, one can find a potentially larger probability space $(Y, \mathscr{D}, \nu)$, commuting actions S_1 and S_2 of G on $(Y, \mathscr{D}, \nu)$, and a measurable, measure-preserving map $\pi : Y \to X$ intertwining T_i and S_i such that

$$\int f_0 \cdot T_1^g f_1 \cdot T_1^g T_2^g f_2\,\mathrm{d}\mu - \int \mathbb{E}(f_0 \circ \pi|\mathsf{A}_1 \vee \mathsf{A}_{12}) \cdot T_1^g \mathbb{E}(f_1 \circ \pi|\mathsf{A}_1 \vee \mathsf{A}_2) \cdot T_1^g T_2^g \mathbb{E}(f_2 \circ \pi|\mathsf{A}_{12} \vee \mathsf{A}_2)\,\mathrm{d}\nu \tag{1.6}$$

averaged along any Følner sequence in G converges to 0, where $\mathsf{A}_1, \mathsf{A}_2$ and A_{12} are the sub-σ-algebras of $\mathscr{D}$ generated by the S_1, S_2 and S_1S_2-invariant functions in $\mathrm{L}^2(Y, \mathscr{D}, \nu)$ respectively. The σ-algebras $\mathsf{A}_1 \vee \mathsf{A}_{12}$, $\mathsf{A}_1 \vee \mathsf{A}_2$ and $\mathsf{A}_{12} \vee \mathsf{A}_2$ are called *characteristic factors* for the correlation (1.5). Austin also gave similar results for longer correlations.

When G is non-amenable, averaging along Følner sequences is unavailable. Recently limits along minimal idempotent ultrafilters (idempotents in the Stone–Čech compactification βG of G that belong to a minimal ideal – see Section 4 for details and [BM07] for the relative merits of minimal idempotents) have been employed as a replacement. It was shown in [BM07] that for any minimal idempotent ultrafilter p on a countable group G one has

$$\lim_{g\to\mathsf{p}} \int f_0 \cdot T_1^g f_1 \cdot T_1^g T_2^g f_2 \,\mathrm{d}\mu - \int f_0 \cdot T_1^g \mathbb{E}(f_1|\mathsf{C}_1) \cdot T_1^g T_2^g \mathbb{E}(f_2|\mathsf{C}_{12}) \,\mathrm{d}\mu = 0 \tag{1.7}$$

where C_1 and C_{12} are the sub-σ-algebras corresponding to functions that are almost-periodic for T_1 and T_1T_2 respectively over A_2. When G is amenable, one can obtain stronger combinatorial results by using minimal idempotent ultrafilters rather than Følner sequences: this is because positivity of correlations along minimal idempotent ultrafilters yields a larger set of $g \in G$ for which (1.1) is positive; see [BM07] for details. For limits of longer correlations along minimal idempotent ultrafilters there is no known description of sub-σ-algebras for which an analogue of (1.7) holds. We remark that the difficulty in adapting the techniques in either [Aus13] or [BM07] lies in the apparent need to understand certain measures that are not invariant, but merely asymptotically invariant along the ultrafilter.

In this paper we combine Austin's satedness techniques with limits along minimal idempotent ultrafilters to obtain the expected characteristic factors for commuting actions of minimally almost periodic groups.

Theorem 1.8. *Let G be a minimally almost periodic group and let T_1, T_2 be commuting, measure-preserving actions of G on a compact, Hausdorff probability space (X, μ) via homeomorphisms. For any $\varepsilon > 0$ and any f_1 in $\mathrm{L}^\infty(X, \mu)$ bounded by 1 there are commuting, measure-preserving actions S_1, S_2 of G on a compact, Hausdorff probability space (Y, ν) and an intertwining factor map $\pi : Y \to X$ such that*

$$\left| \lim_{g\to\mathsf{p}} \int f_0 \cdot T_1^g f_1 \cdot T_1^g T_2^g f_2 \,\mathrm{d}\mu - \int \mathbb{E}(f_0 \circ \pi|\mathsf{A}_1 \vee \mathsf{A}_{12}) \cdot S_1^g \mathbb{E}(f_1 \circ \pi|\mathsf{A}_1 \vee \mathsf{A}_2) \cdot S_1^g S_2^g \mathbb{E}(f_2 \circ \pi|\mathsf{A}_{12} \vee \mathsf{A}_2) \,\mathrm{d}\nu \right| < \varepsilon$$

for all minimal idempotent ultrafilters p on G and all f_0, f_2 in $\mathrm{L}^\infty(X, \mu)$ bounded by 1.

Note that the space X in Theorem 1.8 is not assumed to be metrizable, and that the group G can be uncountable. For this reason we need to extend Austin's notion of satedness to such spaces; this generalization is carried out in Section 3. We then combine Theorem 1.8 with an application of Gelfand theory to obtain the following strong recurrence result.

Theorem 1.9. *Let G be a minimally almost periodic group and let T_1 and T_2 be commuting actions of G on a probability space $(X, \mathscr{B}, \mu)$ by measurable, measure-preserving maps. If μ is ergodic for the $G \times G$ action $(g_1, g_2) \mapsto T_1^{g_1} T_2^{g_2}$ then*

$$\lim_{g\to\mathsf{p}} \int f_0 \cdot T_1^g f_1 \cdot T_1^g T_2^g f_2 \,\mathrm{d}\mu \geq \left(\int f_0^{1/4} f_1^{1/4} f_2^{1/4} \,\mathrm{d}\mu \right)^4 \tag{1.10}$$

for any minimal idempotent ultrafilter p on G and any non-negative measurable functions f_0, f_1, f_2 on X.

We remark that, by [CZK15, Theorem B.1], the exponent in (1.10) cannot be improved to 3 in general.

To deduce the existence of triangles in large enough subsets of quasirandom groups from Theorem 1.9 we need to form a limiting group from a quasirandom sequence $n \mapsto G_n$. For any sequence $n \mapsto G_n$ of finite groups and any ultrafilter on $\mathbb{N}$ one can form their ultraproduct G, which can be given the structure of a probability group using Loeb measure (for details on Loeb measure see, for example [Cut83]). When the sequence $n \mapsto G_n$ is quasirandom the group G is called an *ultra quasirandom group*. There are commuting actions of G on the ultraproduct X of the sequence $n \mapsto G_n \times G_n$ of groups that correspond to left multiplication by G_n in the first and second coordinates of $G_n \times G_n$ respectively. In the case that G is minimally almost periodic, we obtain the following result from Theorem 1.9.

Theorem 1.11. *Let $n \mapsto G_n$ be a sequence of finite groups such that their ultraproduct G is minimally almost periodic, and let Ω be the ultraproduct of the groups $G_n \times G_n$. Let L_1 and L_2 be the actions of G on*

Ω *induced by left multiplication in the first and second coordinates respectively and let* m *be Loeb measure on* Ω*. For any minimal idempotent ultrafilter* p *on* G *we have*

$$\lim_{g\to\mathsf{p}} \int f_0 \cdot L_1^g f_1 \cdot L_1^g L_2^g f_2 \,\mathrm{dm} \geq \left(\int f_0^{1/4} f_1^{1/4} f_2^{1/4} \,\mathrm{dm} \right)^4$$

for any non-negative measurable functions f_0, f_1, f_2 *on* Ω*.*

Theorem 1.11 requires that the ultra quasirandom group determined by $n \mapsto G_n$ is minimally almost periodic. In [BT14] it was shown that any ultraproduct of the sequence $n \mapsto \mathrm{SL}(2, \mathbb{F}_{p^n})$ is minimally almost periodic. More recently, work by Yang [Yan16] provides many examples of classes $\mathcal{F}$ of groups with the property that the ultraproduct of any quasirandom sequence $n \mapsto G_n$ in $\mathcal{F}$ is minimally almost periodic. Such classes are called *q.u.p. (quasirandom ultra product)* classes. For example, the class of finite, quasisimple groups is q.u.p. by [Yan16, Corollary 1.12].

Theorem 1.12. *Let* $\mathcal{F}$ *be a q.u.p. class of finite groups. For every* $0 < \alpha < 1$ *and every* $\varepsilon > 0$ *there exist* $D, K \in \mathbb{N}$ *such that for every* D*-quasirandom group* $G \in \mathcal{F}$ *and every* $A \subset G \times G$ *with* $|A| \geq \alpha |G|^2$ *the set*

$$\left\{ g \in G : \frac{|A \cap (1,g)^{-1}A \cap (g,g)^{-1}A|}{|G|^2} > \alpha^4 - \varepsilon \right\} \tag{1.13}$$

has the property that at most K *of its right shifts are needed to cover* G*.*

It would be interesting to obtain a version Theorem 1.12 with explicit description of D and K. Such a proof may also shed light on the question of how large the set (1.13) can be: we conjecture that its density in G should tend to 1 as $D \to \infty$ in analogy with Theorem 1.2, but have been unable to prove this using our techniques.[1] Combined with the fact that the non-cyclic finite simple groups are quasirandom in the sense that the minimal dimension of a non-trivial irreducible representation grows with the order of the group (see [Gow08, Theorem 4.7]), Theorem 1.12 yields the following consequence:

Corollary 1.14. *For every* $0 < \alpha < 1$ *and every* $\varepsilon > 0$ *there exist* $N, K \in \mathbb{N}$ *such that for every non-cyclic finite simple group* G *of order at least* N *and every* $A \subset G \times G$ *with* $|A| \geq \alpha |G|^2$ *the set* (1.13) *has the property that at most* K *of its right shifts are needed to cover* G*.*

We conclude by mentioning that an alternative method for forming a limiting group from a quasirandom sequence is available when the sequence is increasing. Combinatorially, it gives many $g \in A$ for which $A \cap (1,g)^{-1}A \cap (g,g)^{-1}A \neq \varnothing$, but does not give much information about the size of the intersection. Given a sequence $n \mapsto G_n$ of groups such that $G_n \hookrightarrow G_{n+1}$ for all $n \in \mathbb{N}$, the direct limit G is the union of the embeddings $G_n \hookrightarrow G$ (see [Rob96, Page 23]) and is therefore amenable. If, in addition, the sequence is quasirandom, then G is minimally almost periodic. For such sequences we apply [BCRZK16, Corollary 4.9] to obtain the following combinatorial result.

Theorem 1.15. *Let* $n \mapsto G_n$ *be a quasirandom sequence such that* $G_n \hookrightarrow G_{n+1}$ *for all* $n \in \mathbb{N}$*. For every* $\alpha > 0$ *and every* $\varepsilon > 0$ *there is* $N \in \mathbb{N}$ *such that, for any* $n \geq N$ *and any* $A \subset G_n \times G_n$ *with* $|A| \geq \alpha |G_n|^2$ *there are* $(1-\varepsilon)|G_n|$ *many* $g \in G_n$ *for which* $A \cap (1,g)^{-1}A \cap (g,g)^{-1}A$ *is non-empty.*

The rest of the paper runs as follows. In Section 2 we define the categories of dynamical systems we will work with and recall the Jacobs–de Leeuw–Glicksberg decomposition. A version of satedness suitable for our needs is developed in Section 3. In Sections 4 and 5 we present the necessary facts regarding minimal idempotent ultrafilters and minimally almost periodic groups respectively. Theorem 1.8 is proved in Section 6 and Theorem 1.9 is proved in Section 7. Lastly, the combinatorial results mentioned above are proved in Section 8.

We would like to thank the referees for constructive and thorough reports.

[1] These problems have been solved by Austin [Aus16] after the completion of this article.

2. Topological and Measure-preserving dynamical systems

In this section various categories of dynamical systems are defined that will be used throughout the paper.

Definition 2.1. Let G be a discrete group. The objects of the category $\mathcal{C}^{\text{top}}(G)$ are the left actions $T : G \times X \to X$ of G on compact, Hausdorff spaces X having the property that each of the induced maps $T^g : X \to X$ is continuous. Objects in $\mathcal{C}^{\text{top}}(G)$ will be called *systems* and denoted $\mathbf{X} = (X, T)$. Their defining continuous left actions will be called just *actions*. The morphisms of the category $\mathcal{C}^{\text{top}}$ are the continuous maps intertwining the G actions on the domain and the codomain.

Note that the objects of $\mathcal{C}^{\text{top}}(G)$ may be actions on non-metrizable topological spaces. This level of generality is needed in order to handle actions of very large groups on the Gelfand spaces of non-separable C*-algebras, which will play a role in the next section.

In our applications the acting group will be G^2 and an action in $\mathcal{C}^{\text{top}}(G^2)$ will be written in the form $(g_1, g_2) \mapsto T_1^{g_1} T_2^{g_2}$ where T_1 and T_2 are commuting G actions. We also write $T_{12}^g = T_1^g T_2^g$.

For any system $\mathbf{X}$ we denote by $M_{\mathbf{X}}$ the set of Baire probability measures on X that are T-invariant. (Recall that the *Baire* sets are the members of the σ-algebra generated by the compact G_δ sets.) In view of the Riesz–Markov–Kakutani representation theorem this set can be seen as a subset of the dual of the space $\mathrm{C}(X)$ of continuous, real-valued functions on X equipped with the uniform norm, and as such it is compact and convex with respect to the weak* topology. The set of extreme points of $M_{\mathbf{X}}$ is denoted $\operatorname{ex} M_{\mathbf{X}}$. It follows from the Radon–Nikodym theorem that the extreme points of $M_{\mathbf{X}}$ are precisely the ergodic measures, namely the measures for which every almost invariant Baire set has measure either 0 or 1.

Definition 2.2. The objects of the category $\mathcal{C}^{\text{meas}}$ are pairs $(\mathbf{X}, \mu)$, where $\mu \in M_{\mathbf{X}}$, called *measure-preserving systems*. A morphism $(\mathbf{Y}, \nu) \to (\mathbf{X}, \mu)$ in $\mathcal{C}^{\text{meas}}$ is any morphism $\pi : \mathbf{Y} \to \mathbf{X}$ in $\mathcal{C}^{\text{top}}$ such that $\pi\nu = \mu$. When there is a morphism $\pi : (\mathbf{Y}, \nu) \to (\mathbf{X}, \mu)$ in $\mathcal{C}^{\text{meas}}$ we call π the *factor map* and say that $(\mathbf{Y}, \nu)$ is an *extension* of $(\mathbf{X}, \mu)$, or that $(\mathbf{X}, \mu)$ is a *factor* of $(\mathbf{Y}, \nu)$. The category $\mathcal{C}^{\text{erg}}$ is the subcategory of $\mathcal{C}^{\text{meas}}$ whose objects are the pairs $(\mathbf{X}, \mu)$ for which $\mu \in \operatorname{ex} M_{\mathbf{X}}$. The objects of $\mathcal{C}^{\text{erg}}$ are called *ergodic systems*.

Lemma 2.3. *Let $\psi : \mathbf{Y} \to \mathbf{X}$ be a morphism in $\mathcal{C}^{\text{top}}$ and fix μ in $\operatorname{ex} M_{\mathbf{X}}$. Then $\tilde{M}_{\mathbf{Y}} = \{\nu \in M_{\mathbf{Y}} : \psi\nu = \mu\}$ is a compact, convex set and its extreme points are ergodic measures on $\mathbf{Y}$.*

Proof. It is clearly closed and convex. To see that its extreme points are ergodic, it suffices to show that $\operatorname{ex} \tilde{M}_{\mathbf{Y}} \subset \operatorname{ex} M_{\mathbf{Y}}$. Indeed, suppose $\lambda \in \operatorname{ex} \tilde{M}_{\mathbf{Y}}$ can be written as $\lambda = c\lambda_1 + (1-c)\lambda_2$ with $\lambda_1, \lambda_2 \in M_{\mathbf{Y}}$. Then $c\pi_*\lambda_1 + (1-c)\pi_*\lambda_2 = \mu$, so by extremality of μ in $M_{\mathbf{X}}$ we have $\pi_*\lambda_1 = \pi_*\lambda_2 = \mu$. By extremality of λ in $\tilde{M}_{\mathbf{Y}}$ this implies $\lambda_1 = \lambda_2 = \lambda$. □

We will be concerned with sub-σ-algebras of invariant sets. Given a measure space $(X, \mathscr{B}, \mu)$ and sets $A, B \in \mathscr{B}$, write $A \sim B$ when $\mu(A \triangle B) = 0$.

Definition 2.4. Let $(\mathbf{X}, \mu)$ be a measure-preserving system in $\mathcal{C}^{\text{meas}}(G)$. A Baire subset $B \subset X$ is *almost invariant with respect to* μ if $(T^g)^{-1}B \sim B$ for every $g \in G$. Write $\mathsf{A}^\mu\mathbf{X}$ for the sub-σ-algebra generated by the almost invariant sets. If $(\mathbf{X}, \mu)$ is a measure-preserving system in $\mathcal{C}^{\text{meas}}(G^2)$ and $i \in \{1, 2, 12\}$, write $\mathsf{A}_i^\mu\mathbf{X}$ for the sub-σ-algebra generated by the T_i almost invariant sets. Lastly, for any $i, j \in \{1, 2, 12\}$ define $\mathsf{A}_{i,j}^\mu\mathbf{X} = \mathsf{A}_i^\mu\mathbf{X} \vee \mathsf{A}_j^\mu\mathbf{X}$.

We conclude this section by recalling a general version of the splitting of $\mathrm{L}^2(\mathbf{X}, \mu)$ into almost periodic and weakly mixing parts that will be needed in the proof of Theorem 1.9. Let G be a group and let $(\mathbf{X}, \mu)$ be a measure preserving system in $\mathcal{C}^{\text{meas}}(G)$. Write $\mathscr{H}$ for $\mathrm{L}^2(\mathbf{X}, \mu)$. Consider the collection $\mathscr{S} = \{T^g : g \in G\}$ of unitary operators on $\mathscr{H}$. The closure with respect to the weak topology of any orbit of $\mathscr{S}$ is compact in the weak topology, so $\mathscr{S}$, in the terminology of [LG61], is a weakly almost periodic semigroup of operators. Write $\overline{\mathscr{S}}$ for the closure of $\mathscr{S}$ in the weak operator topology. Applying [LG61, Corollary 4.12] allows us to write $\mathscr{H} = \mathscr{H}_{\mathsf{o}} + \mathscr{H}_{\mathsf{r}}$ where

$$\mathscr{H}_{\mathsf{o}} = \{f \in \mathscr{H} : 0 \in \overline{\mathscr{S}} f\}$$

is the closed subspace of *flight vectors* and

$$\mathscr{H}_{\mathsf{r}} = \{f \in \mathscr{H} : \overline{\mathscr{S}} f = \overline{\mathscr{S}} h \text{ for every } h \in \overline{\mathscr{S}} f\}$$

is the closed subspace of *reversible vectors*. Since $\mathscr{S}$ is a group [LG61, Lemma 4.5] implies that $\mathscr{H}_{\mathsf{r}}$ is spanned by the finite dimensional, $\mathscr{S}$-invariant subspaces of $\mathscr{H}$. The splitting $\mathscr{H}_{\mathsf{o}} + \mathscr{H}_{\mathsf{r}}$ is determined by the unique projection in the kernel of $\overline{\mathscr{S}}$. By [LG61, Theorem 2.3(iv)] this kernel is self-adjoint so the unique projection in the kernel is also self-adjoint, implying that the above splitting is orthogonal.

Let $(\mathbf{X}_1, \mu_1)$ and $(\mathbf{X}_2, \mu_2)$ be measure-preserving systems in $\mathcal{C}^{\mathrm{meas}}(G)$ and let $\mathscr{H}_1$ and $\mathscr{H}_2$ be the corresponding Hilbert spaces. The Hilbert space corresponding to the product system $(\mathbf{X}_1 \times \mathbf{X}_2, \mu_1 \times \mu_2)$ is $\mathscr{H}_1 \otimes \mathscr{H}_2$. We will need the following result, which relates the splittings of these Hilbert spaces, in the proof of Theorem 1.9.

Lemma 2.5. *Let $(\mathbf{X}_1, \mu_1)$ and $(\mathbf{X}_2, \mu_2)$ be measure-preserving systems in $\mathcal{C}^{\mathrm{meas}}(G)$ and let $\mathscr{H}_1$ and $\mathscr{H}_2$ be the associated L^2 spaces. Let $\mathscr{H}$ be the L^2 space of the product measure-preserving system $(\mathbf{X}_1 \times \mathbf{X}_2, \mu_1 \otimes \mu_2)$. Then $\mathscr{H}_{1,\mathsf{r}} \otimes \mathscr{H}_{2,\mathsf{r}} = \mathscr{H}_{\mathsf{r}}$.*

Proof. Write

$$\mathscr{H}_{\mathsf{o}} \oplus \mathscr{H}_{\mathsf{r}} = \mathscr{H} = (\mathscr{H}_{1,\mathsf{r}} \otimes \mathscr{H}_{2,\mathsf{r}}) \oplus (\mathscr{H}_{1,\mathsf{o}} \otimes \mathscr{H}_{2,\mathsf{r}}) \oplus (\mathscr{H}_{1,\mathsf{r}} \otimes \mathscr{H}_{2,\mathsf{o}}) \oplus (\mathscr{H}_{1,\mathsf{o}} \otimes \mathscr{H}_{2,\mathsf{o}})$$

and note that $\mathscr{H}_{1,\mathsf{r}} \otimes \mathscr{H}_{2,\mathsf{r}} \subset \mathscr{H}_{\mathsf{r}}$, whereas if $f \in \mathscr{H}_{1,\mathsf{o}}$ and $g \in \mathscr{H}_2$, or if $f \in \mathscr{H}_1$ and $g \in \mathscr{H}_{2,\mathsf{o}}$, then $f \otimes g \in \mathscr{H}_{\mathsf{o}}$. □

3. Satedness

Satedness, introduced by Austin [Aus10; Aus13] to prove convergence of multiple ergodic averages, is a property that a measure-preserving system may possess with respect to certain classes of systems. For example, one could speak of satedness with respect to the class of Kronecker systems. Although the concept depends critically on an invariant measure, it will be convenient to consider classes defined by topological, rather than measure-theoretic properties. We will therefore consider satedness with respect to idempotent classes, defined below, in $\mathcal{C}^{\mathrm{top}}$. Recall that a *joining* of objects $\mathbf{X}_1, \dots, \mathbf{X}_n$ in $\mathcal{C}^{\mathrm{top}}$ is an object $\mathbf{Z}$ in $\mathcal{C}^{\mathrm{top}}$ together with factor maps $\mathbf{Z} \to \mathbf{X}_i$ for each $1 \le i \le n$.

Definition 3.1. An *idempotent class* in $\mathcal{C}^{\mathrm{top}}$ is a class that contains the one-point system, is closed under joinings of finitely many systems, and is closed under arbitrary inverse limits.

It follows from Zorn's lemma that every idempotent class $\mathcal{I}$ defines a map $\mathcal{I} : \mathbf{X} \mapsto \mathcal{I}\mathbf{X}$ on $\mathcal{C}^{\mathrm{top}}$ that associates to every system its maximal factor in $\mathcal{I}$. We remark that this map can be seen as a natural transformation from the identity functor on $\mathcal{C}^{\mathrm{top}}$.

The following idempotent classes on $\mathcal{C}^{\mathrm{top}}(G^2)$ will play a crucial role in what follows.

Definition 3.2. Let G be a group and let $\mathbf{X}$ be a system in $\mathcal{C}^{\mathrm{top}}(G^2)$. For $i = 1, 2, 12$ let $\mathcal{I}_i$ be the idempotent class of systems on which the action T_i is trivial. Define $\mathcal{I}_{i,j} = \mathcal{I}_i \vee \mathcal{I}_j$ for any $i, j \in \{1, 2, 12\}$.

Note that, for any $\mu \in M_{\mathbf{X}}$ the space $\mathrm{L}^2(\mathcal{I}_{1,2}\mathbf{X}, \mu)$ is a priori smaller than $\mathrm{L}^2(\mathsf{A}^{\mu}_{1,2}\mathbf{X}, \mu)$.

We now turn to our version of satedness, which is based on Austin's definition in [Aus13], but differs in that the energy increment (3.4) below is quantified, rather than being required to vanish.

Definition 3.3. Let $\mathcal{C}$ be a subcategory of $\mathcal{C}^{\mathrm{meas}}$ and let $\mathcal{I}$ be an idempotent class (or, more generally, a natural transformation from the identity functor in $\mathcal{C}^{\mathrm{top}}$). A measure-preserving system $(\mathbf{X}, \mu)$ is called $(\varepsilon, f, \mathcal{I})$ *sated in* $\mathcal{C}$ for $\varepsilon \ge 0$ and $f \in \mathrm{L}^2(\mathbf{X}, \mu)$ if one of the following equivalent conditions holds.

(1) For every extension $\pi : (\mathbf{Y}, \nu) \to (\mathbf{X}, \mu)$ in $\mathcal{C}$ we have

$$\|\mathbb{E}_\nu(f \circ \pi | \mathcal{I}\mathbf{Y}) - \mathbb{E}_\mu(f | \mathcal{I}\mathbf{X}) \circ \pi\|_2 \le \varepsilon. \tag{3.4}$$

(2) For every extension $\pi : (\mathbf{Y}, \nu) \to (\mathbf{X}, \mu)$ in $\mathcal{C}$ and every $\phi \in \mathrm{C}(\mathcal{I}\mathbf{Y})$ we have

$$\left| \int f \circ \pi \cdot \phi \, \mathrm{d}\nu - \int \mathbb{E}_\mu(f | \mathcal{I}\mathbf{X}) \circ \pi \cdot \phi \, \mathrm{d}\nu \right| \le \varepsilon \|\phi\|_2. \tag{3.5}$$

The conditions (3.4) and (3.5) are equivalent because $\|v\| = \sup\{|\langle v, w\rangle| : \|w\| = 1\}$ in any Hilbert space and $\mathrm{C}(\mathcal{I}\mathbf{Y})$ is dense in $\mathrm{L}^2(\mathcal{I}\mathbf{Y}, \mu)$.

The following result shows that every measure-preserving system has an extension that is sated up to any prescribed error. We will be able to do this without the use of inverse limits because the energy increment

in (3.4) is only required to be small rather than to vanish. This extends the applicability of satedness to categories where inverse limits may not exist.

Theorem 3.6. *Let $\mathcal{C}$ be a subcategory of $\mathcal{C}^{\mathrm{meas}}$ and let $\mathcal{I}$ be an idempotent class in $\mathcal{C}^{\mathrm{top}}$, or more generally a natural transformation from the identity functor in $\mathcal{C}^{\mathrm{top}}$. Let $(\mathbf{X},\mu)$ be a measure-preserving system in the category $\mathcal{C}$. Then for any $\varepsilon>0$ and any $f\in \mathrm{L}^2(\mathbf{X},\mu)$ there exists an extension $\psi:(\mathbf{Y},\nu)\to(\mathbf{X},\mu)$ of measure-preserving systems such that $(\mathbf{Y},\nu)$ is $(\varepsilon, f\circ\psi,\mathcal{I})$ sated in $\mathcal{C}$.*

Proof. Fix $\varepsilon>0$ and $f\in \mathrm{L}^2(\mathbf{X},\mu)$. Assume $f\neq 0$ as otherwise the conclusion is immediate. We have

$$\sup\{\|\mathbb{E}(f\circ\psi|\mathcal{I}\mathbf{Y})\|_2 : (\mathbf{Y},\nu)\xrightarrow{\psi}(\mathbf{X},\mu) \text{ a morphism}\}\leq \|f\|_2<\infty,$$

so there exists an extension $\psi:(\mathbf{Y},\nu)\to(\mathbf{X},\mu)$ such that $\|\mathbb{E}(f\circ\psi|\mathcal{I}\mathbf{Y})\|$ is within $\varepsilon/(2\|f\|_2)$ of the supremum above. For any further extension $\pi:(\mathbf{Z},\lambda)\to(\mathbf{Y},\nu)$ we have

$$\mathbb{E}(f\circ\psi|\mathcal{I}\mathbf{Y})\circ\pi=\mathbb{E}(f\circ\psi\circ\pi|\pi^{-1}(\mathcal{I}\mathbf{Y}))=\mathbb{E}(\mathbb{E}(f\circ\psi\circ\pi|\mathcal{I}\mathbf{Z})|\pi^{-1}(\mathcal{I}\mathbf{Y}))$$

by functoriality of $\mathcal{I}$. On the other hand, by choice of $\mathbf{Y}$ we have

$$\|\mathbb{E}(f\circ\psi|\mathcal{I}\mathbf{Y})\|_2>\|\mathbb{E}(f\circ\psi\circ\pi|\mathcal{I}\mathbf{Z})\|_2-\delta.$$

Since the conditional expectation onto $\pi^{-1}(\mathcal{I}\mathbf{Y})$ is an orthogonal projection, this implies

$$\|\mathbb{E}(f\circ\psi|\mathcal{I}\mathbf{Y})-\mathbb{E}(f\circ\psi\circ\pi|\mathcal{I}\mathbf{Z})\|_2\leq 2\|\mathbb{E}(f\circ\psi\circ\pi|\mathcal{I}\mathbf{Z})\|_2\delta\leq 2\|f\|_2\delta,$$

so $(\mathbf{Y},\nu)$ is $(\varepsilon,f\circ\psi,\mathcal{I})$ sated provided $\delta<\varepsilon/(2\|f\|_2)$. □

Theorem 3.6 will be applied to the category $\mathcal{C}^{\mathrm{erg}}$ of ergodic systems. However, we will need satedness in the class of *all* measure-preserving systems. Switching between these classes requires a version of the ergodic decomposition. Since we do not assume metrizability of the compact spaces under consideration, we use the following Choquet-type theorem. Recall that a function f from a convex set M to $\mathbb{R}$ is *affine* if $f(tx+(1-t)y)=tf(x)+(1-t)f(y)$ for all $0\leq t\leq 1$ and all $x,y\in M$.

Theorem 3.7 (Choquet–Bishop–de Leeuw, [Phe01, p. 17])**.** *Suppose that M is a compact convex subset of a locally convex space and let $\mu\in M$. Then there exists a probability measure η on M that represents μ in the sense that*

$$\phi(\mu)=\int \phi\,\mathrm{d}\eta$$

for every continuous affine function $\phi:M\to\mathbb{R}$ and such that η vanishes on every Baire subset of M that is disjoint from $\operatorname{ex} M$.

In this version of the Choquet theorem the representing measure η is not unique and is only supported by the extreme points in a weak sense, but this will not be an issue.

Lemma 3.8. *Let $(\mathbf{X},\mu)$ be a measure-preserving system and $F\subset \mathrm{L}^\infty(X)$. Then there exists an extension $\pi:(\mathbf{Y},\nu)\to(\mathbf{X},\mu)$ such that $f\circ\pi$ coincides with a continuous function on Y ν-a.e. for every $f\in F$.*

Proof. Let A be the minimal G-invariant C^*-subalgebra of $\mathrm{L}^\infty(X,\mu)$ that contains $\mathrm{C}(X)\cup F$. Let Y be its Gelfand spectrum with the canonical G-action and the canonical projection π onto X. We have a positive linear functional ν on A given by $\nu(g)=\int g\,\mathrm{d}\mu$, this defines a G-invariant probability measure on Y.

It remains to show that $f\circ\pi$ coincides with a continuous function ν-a.e. for every $f\in F$. Fix $f\in F$, by duality it suffices to verify

$$\int (f\circ\pi)g\,\mathrm{d}\nu=\int \tilde{f}g\,\mathrm{d}\nu \tag{3.9}$$

for every $g\in L^1(Y,\nu)$, where $\tilde{f}$ is the continuous function on Y corresponding to f viewed as an element of A. Both $f\circ\pi$ and $\tilde{f}$ are bounded functions, so it suffices to verify this identity for g in a dense subspace of $L^1(Y,\nu)$. We claim that $\pi^*\,\mathrm{C}(X)$ is one such subspace. Indeed, $\mathrm{C}(Y)$ is dense in $L^1(Y,\nu)$ and $\mathrm{C}(X)$ is L^1-dense in A. This implies that $\pi^*\,\mathrm{C}(X)$ is L^1 dense in $\mathrm{C}(Y)$. For every $g\in\mathrm{C}(X)$ we have $\tilde{g}=g\circ\pi$, so (3.9) boils down to

$$\int ((fg)\circ\pi)\,\mathrm{d}\nu=\int fg\,\mathrm{d}\mu$$

as desired. □

Proposition 3.10. *Let $(\mathbf{X},\mu)$ be $(\varepsilon,f,\mathcal{I})$-sated in $\mathcal{C}^{\mathrm{erg}}$ for some $f\in \mathrm{L}^\infty(X,\mu)$ and $\varepsilon\geq 0$. Then $(\mathbf{X},\mu)$ is also $(\varepsilon,f,\mathcal{I})$-sated in $\mathcal{C}^{\mathrm{meas}}$.*

Proof. Let $\pi:(\mathbf{Y},\nu)\to(\mathbf{X},\mu)$ be an extension in $\mathcal{C}^{\mathrm{meas}}$. We have to show (3.5) for every $\phi\in \mathrm{C}(\mathcal{I}\mathbf{Y})$. Passing to a further extension of $(\mathbf{Y},\nu)$ using Lemma 3.8, we may assume that both f and $\tilde{f}=\mathbb{E}_\mu(f|\mathcal{I}X)\circ\pi$ (which are a priori merely bounded measurable functions) admit representatives in $\mathrm{C}(Y)$. While verifying

$$\left|\int f\phi\,\mathrm{d}\nu-\int \tilde{f}\phi\,\mathrm{d}\nu\right|\leq \varepsilon\|\phi\|_{\mathrm{L}^2(\nu)}$$

for every $\phi\in\mathrm{C}(\mathcal{I}\mathbf{Y})$ there is no harm in replacing f and $\tilde{f}$ by their continuous representatives.

Write $\tilde{M}_{\mathbf{Y}}$ for the measures in $M_{\mathbf{Y}}$ that extend μ. By Lemma 2.3 it is a closed, convex subset of $M_{\mathbf{Y}}$ whose extreme points are ergodic. Thus for any measure $\lambda\in\operatorname{ex}\tilde{M}_{\mathbf{Y}}$ we have

$$\left|\int f\phi\,\mathrm{d}\lambda-\int \tilde{f}\phi\,\mathrm{d}\lambda\right|\leq \varepsilon\|\phi\|_{\mathrm{L}^2(\lambda)}. \tag{3.11}$$

by the satedness hypothesis. Let now η be a measure on $\tilde{M}_{\mathbf{Y}}$ representing ν in the sense of the Choquet–Bishop–de Leeuw theorem (Theorem 3.7). Consider the set

$$\Lambda=\left\{\lambda\in\tilde{M}_{\mathbf{Y}}:\left|\int f\phi\,\mathrm{d}\lambda-\int \tilde{f}\phi\,\mathrm{d}\lambda\right|>\varepsilon\|\phi\|_{\mathrm{L}^2(\lambda)}\right\}$$

which is disjoint from $\operatorname{ex}\tilde{M}_{\mathbf{Y}}$ in view of (3.11), and Baire because it consists of those measures λ where one continuous function of λ is larger than another. It follows that $\eta(\Lambda)=0$. Therefore

$$\begin{aligned}\left|\int f\phi\,\mathrm{d}\nu-\int \tilde{f}\phi\,\mathrm{d}\nu\right|&\leq \int\limits_{\tilde{M}_{\mathbf{Y}}}\left|\int f\phi\,\mathrm{d}\lambda-\int \tilde{f}\phi\,\mathrm{d}\lambda\right|\mathrm{d}\eta(\lambda)\\&=\int\limits_{\tilde{M}_{\mathbf{Y}}\setminus\Lambda}\left|\int f\phi\,\mathrm{d}\lambda-\int \tilde{f}\phi\,\mathrm{d}\lambda\right|\mathrm{d}\eta(\lambda)\leq\int\limits_{\tilde{M}_{\mathbf{Y}}\setminus\Lambda}\varepsilon\|\phi\|_{L^2(\lambda)}\,\mathrm{d}\eta(\lambda)\leq\varepsilon\|\phi\|_{\mathrm{L}^2(\nu)}\end{aligned}$$

by Hölder's inequality. □

4. Minimal idempotent ultrafilters

For any non-empty set X write βX for the collection of ultrafilters on X. Recall that these are the filters on X that are maximal with respect to containment, and can be thought of as finitely-additive $\{0,1\}$-valued measures on X. We identify each $x\in X$ with the principal ultrafilter $\delta_x=\{A\subset X:x\in A\}$. Upon equipping βX with the topology defined by the base consisting of the clopen sets $\overline{A}=\{\mathsf{p}\in\beta X:A\in\mathsf{p}\}$ for any subset A of X it becomes a compact, Hausdorff topological space. It enjoys the following universal property, which will be used repeatedly as a means to take limits along ultrafilters.

Proposition 4.1. *Let X be a non-empty set. For any compact, Hausdorff topological space Z and any map $\phi:X\to Z$ there is a continuous map $\beta X\to Z$ that agrees with ϕ on the principal ultrafilters.*

Given a map ϕ from X to a compact, Hausdorff space, we denote by

$$\lim_{x\to\mathsf{p}}\phi(x)$$

the value at $\mathsf{p}\in\beta X$ of the extension provided by Proposition 4.1.

Let G be any group. One can make βG a semigroup by defining

$$\mathsf{p}*\mathsf{q}=\{A\subset G:\{g\in G:Ag^{-1}\in\mathsf{p}\}\in\mathsf{q}\} \tag{4.2}$$

for any $\mathsf{p},\mathsf{q}\in\beta G$. Note that $\delta_g*\delta_h=\delta_{gh}$ for any $g,h\in G$ so (4.2) extends multiplication on G. The operation above makes βG a right semi-topological semigroup: for any fixed p in βG the map $\mathsf{q}\mapsto\mathsf{p}*\mathsf{q}$ is continuous. Ellis's lemma [Ell58, Lemma 1] implies that there are idempotents for (4.2) in any compact sub-semigroup of βG. The semigroup operation on βG interacts with continuous actions of G on compact, Hausdorff spaces in the following way (cf. [BH90, Lemma 6.1]).

Proposition 4.3. *Let G be a group and let $T : G \times X \to X$ be a right actions of G on a compact, Hausdorff space X via continuous maps. Then*

$$\lim_{g \to \mathsf{p} * \mathsf{q}} T^g x = \lim_{g \to \mathsf{q}} \lim_{h \to \mathsf{p}} T^g (T^h x)$$

for all $x \in X$ and all $\mathsf{p}, \mathsf{q} \in \beta G$. In particular, if p is idempotent then

$$\lim_{g \to \mathsf{p}} T^g x = \lim_{g \to \mathsf{p}} \lim_{h \to \mathsf{p}} T^g (T^h x)$$

for all $x \in X$.

The following version of the van der Corput trick follows immediately from the proof of [Sch07, Lemma 4].

Proposition 4.4. *Let G be a group and let $\mathscr{H}$ be a Hilbert space. For any sequence $u : G \to \mathscr{H}$ that is norm-bounded and any idempotent ultrafilter p on G, if*

$$\left| \lim_{h \to \mathsf{p}} \lim_{g \to \mathsf{p}} \langle u(hg), u(g) \rangle \right| < \varepsilon$$

then $\| \lim_{g \to \mathsf{p}} u(g) \|^2 \leq \varepsilon$, the latter limit being taken in the weak topology on $\mathscr{H}$.

An idempotent ultrafilter p is *minimal* if it belongs to a minimal right ideal in βG. Every non-zero right ideal contains a minimal right ideal, so Ellis's lemma [Ell58, Lemma 1] implies that every right ideal contains a minimal idempotent ultrafilter. See e.g. [BM07, Section 2] for the details. The following lemma tells us that sets in minimal idempotent ultrafilters have a certain largeness property. Recall that $S \subset G$ is *right syndetic* if there is a finite subset F of G for which $SF = G$.

Lemma 4.5 (cf. [Ber03, Theorem 2.4])**.** *Let G be a group and let p be a minimal idempotent ultrafilter on G. For any $A \in \mathsf{p}$ the set $A^{-1}A = \{g^{-1}h : g, h \in A\}$ is right syndetic.*

Proof. Fix $A \in \mathsf{p}$. Let X be a minimal right ideal containing p. Since $\mathsf{p} * \beta G$ is a right ideal contained in X it must be equal to X, so continuity of the map $\mathsf{q} \mapsto \mathsf{p} * \mathsf{q}$ implies X is compact. Consider the continuous right action T of G on X defined by $T^g(\mathsf{p}) = \mathsf{p} * \delta_g$. The set $U := \overline{A} \cap X$ is open in X and contains p. We claim that the collection $\{(T^g)^{-1}U : g \in G\}$ covers X. Indeed, if not then the complement V of its union is a closed, non-empty, T-invariant subset of X. This implies that V is a right ideal, because any q in V satisfies $\mathsf{q} * \beta G = \mathsf{q} * \overline{G} = \mathrm{cl}(\mathsf{q} * G) \subset \mathrm{cl}(V) = V$ where cl denotes the closure of a set in βG.

Since the sets $(T^g)^{-1}U$ cover X we can extract a finite subcover $(T^{g_1})^{-1}U, \ldots, (T^{g_n})^{-1}U$. Thus for every $g \in G$ there is some $1 \leq i \leq n$ for which $T^g(\mathsf{p}) \in (T^{g_i})^{-1}U$. We can rewrite this as $\mathsf{p} * \delta_{gg_i} \in U$, which is the same as $A(gg_i)^{-1} \in \mathsf{p}$. Putting $F = \{g_1, \ldots, g_n\}$, we have proved that $\{g \in G : Ag^{-1} \in \mathsf{p}\}$ is right syndetic. Since $A \in \mathsf{p}$ the larger set $\{g \in G : Ag^{-1} \cap A \in \mathsf{p}\}$ is also right syndetic. But $g \in A^{-1}A$ if and only if $Ag^{-1} \cap A$ is non-empty, so $A^{-1}A \supset \{g \in G : Ag^{-1} \cap A \in \mathsf{p}\}$ is right syndetic, as desired. □

A subset of a group G is a *right central* set if it belongs to some minimal idempotent ultrafilter, and a *right central** set if it belongs to every minimal idempotent ultrafilter.

Lemma 4.6. *Let G be a group and let $A \subset G$ be right central* set. Then A is right syndetic.*

Proof. Suppose A is not right syndetic. Then its complement B is right thick, meaning that for every finite subset F of G there is some $h \in G$ such that $hF \subset B$. This implies that $\mathscr{F} = \{Bg^{-1} : g \in G\}$ is a filter, so there are ultrafilters on G containing $\mathscr{F}$. The collection I of ultrafilters that contain $\mathscr{F}$ is a closed subset of βG. Moreover, it is a right-ideal, for if $\mathsf{p} \supset \mathscr{F}$ and $\mathsf{q} \in \beta G$ then $\mathsf{p} * \mathsf{q}$ contains $\mathscr{F}$ by (4.2). As remarked above, any right ideal contains a minimal right ideal, so there is a minimal idempotent in I. This implies that B is right central, so A is not right central*. □

5. Minimally almost periodic groups

Let G be any group. Denote by $\mathrm{C_b}(G)$ the Banach space of all bounded functions $f : G \to \mathbb{C}$ equipped with the supremum norm. The G actions L and R on $\mathrm{C_b}(G)$, defined by $(\mathsf{L}_g f)(x) = f(gx)$ and $(\mathsf{R}_g f)(x) = f(xg)$ respectively, are isometric. A function $f \in \mathrm{C_b}(G)$ is called *almost periodic* if the subset $\{\mathsf{L}_g f : g \in G\}$ of $\mathrm{C_b}(G)$ is relatively compact. Given a representation ϕ of G on a finite-dimensional, complex Hilbert space

V and vectors x, y in V, the function $f(g) = \langle \phi(g)x, y\rangle$ is almost-periodic. A group G is *minimally almost periodic* if the only almost periodic functions on G are the constant functions.

The following result, a version of [BM07, Theorem 2.2], will be used repeatedly below.

Theorem 5.1. *Let G be a minimally almost periodic group, let $(\mathbf{X}, \mu)$ be a measure-preserving system, and let p be a minimal idempotent ultrafilter on G. For any f in $\mathrm{L}^2(\mathbf{X}, \mu)$ we have*

$$\lim_{g\to\mathsf{p}} T^g f = \mathbb{E}(f|\mathsf{A}^\mu\mathbf{X}) \tag{5.2}$$

in the weak topology of $\mathrm{L}^2(\mathbf{X}, \mu)$.

Proof. Fix $f \in \mathrm{L}^2(\mathbf{X}, \mu)$. Equipped with the weak topology, the unit ball of $\mathrm{L}^2(\mathbf{X}, \mu)$ is compact and Hausdorff so the limit in (5.2) makes sense via Proposition 4.1. Let ϕ be the limit of the sequence $T^g f$ along p. We first show that ϕ belongs to $\mathrm{L}^2(X, \mathsf{A}^\mu\mathbf{X}, \mu)$.

We claim that the orbit $\{T^g\phi : g \in G\}$ is relatively compact in the norm topology. Fix $\varepsilon > 0$. We have

$$\lim_{h\to\mathsf{p}} T^h\phi = \lim_{h\to\mathsf{p}} T^h \lim_{g\to\mathsf{p}} T^g f = \lim_{g\to\mathsf{p}*\mathsf{p}} T^g f = \phi$$

by Proposition 4.3 because p is idempotent. Combined with

$$\|T^g\phi - \phi\|^2 = \langle T^g\phi, T^g\phi\rangle - \langle T^g\phi, \phi\rangle - \langle \phi, T^g\phi\rangle + \langle \phi, \phi\rangle$$

we see that $A := \{g \in G : \|T^g\phi - \phi\| < \varepsilon/2\} = A^{-1}$ belongs to p. Thus AA^{-1} is syndetic by Lemma 4.5. Let $F \subset G$ be finite with $AA^{-1}F = G$. Fix $g \in G$ and write $g = ab^{-1}k$ accordingly. We see that

$$\|T^g\phi - T^k\phi\| = \|T^{ab^{-1}}\phi - \phi\| = \|T^a\phi - T^b\phi\| \leq \varepsilon$$

so the orbit $\{T^g\phi : g \in G\}$ is covered by the balls of radius ε centered at $T^k\phi$ as k runs through F.

It follows that for any $\xi \in \mathrm{L}^2(\mathbf{X}, \mu)$ the function $g \mapsto \langle T^g\phi, \xi\rangle$ is almost periodic. It is therefore constant because G is minimally almost periodic. Thus $T^g\phi = \phi$ in $\mathrm{L}^2(\mathbf{X}, \mu)$ for every $g \in G$. Let φ be a representative of ϕ. We have $T^g\varphi \sim \varphi$ for every $g \in G$, where $\sim$ denotes equality almost everywhere. Since $\mathsf{A}^\mu\mathbf{X}$ contains the measure zero sets, it follows that φ is $\mathsf{A}^\mu\mathbf{X}$ measurable and that $\phi \in \mathrm{L}^2(X, \mathsf{A}^\mu\mathbf{X}, \mu)$.

Lastly, for any $\psi \in \mathrm{L}^2(X, \mathsf{A}^\mu\mathbf{X}, \mu)$ we have

$$\int \phi\cdot\psi\,\mathrm{d}\mu = \lim_{g\to\mathsf{p}}\int T^g f\cdot\psi\,\mathrm{d}\mu = \lim_{g\to\mathsf{p}}\int f\cdot(T^g)^{-1}\psi\,\mathrm{d}\mu = \int f\cdot\psi\,\mathrm{d}\mu$$

so ϕ is the orthogonal projection of f on $\mathrm{L}^2(X, \mathsf{A}^\mu\mathbf{X}, \mu)$. □

6. Characteristic factors in sated systems

In this section we prove Theorem 1.8. To do so we need the following construction of a relatively independent self-joining of a measure-preserving system $(\mathbf{X}, \mu)$ in $\mathcal{C}^{\mathrm{meas}}(G^2)$ over $\mathsf{A}_2^\mu\mathbf{X}$, by which we mean a measure ν on $\mathbf{X}\times\mathbf{X}$ satisfying (6.2) below. Usually (see [Fur81, Chapter 5], for example) one would construct such a joining using a disintegration of μ over $\mathsf{A}_2^\mu\mathbf{X}$, but the existence of such a disintegration is not clear when X is non-metrizable. In our setting, the need for such a disintegration can be circumvented by using limits along minimal idempotent ultrafilters to give an explicit description of the ergodic projection.

Lemma 6.1. *Let G be a minimally almost periodic group and let $(\mathbf{X}, \mu)$ be a measure-preserving system in $\mathcal{C}^{\mathrm{meas}}(G)$. Then there exists a unique Baire measure ν on $X\times X$ such that*

$$\int f_1\otimes f_2\,\mathrm{d}\nu = \int \mathbb{E}(f_1|\mathsf{A}^\mu\mathbf{X})\cdot\mathbb{E}(f_2|\mathsf{A}^\mu\mathbf{X})\,\mathrm{d}\mu \tag{6.2}$$

for any $f_1, f_2 \in \mathrm{C}(X)$.

Proof. Uniqueness follows immediately by density of $\mathrm{C}(X)\otimes\mathrm{C}(X)$ in $\mathrm{C}(X^2)$, so it remains to show the existence. To this end fix a minimal idempotent ultrafilter p on G. Since G is minimally almost periodic Theorem 5.1 implies that

$$\lim_{g\to\mathsf{p}} T_2^g f = \mathbb{E}(f|\mathsf{A}^\mu\mathbf{X})$$

in the weak topology of $\mathrm{L}^2(\mathbf{X}, \mu)$ for every $f \in \mathrm{L}^2(\mathbf{X}, \mu)$.

Let $\delta : X \to X^2$ be the diagonal embedding and let λ be the push-forward $\delta\mu$. Define an action R of G on X^2 by $R^g(x_1, x_2) = (x_1, T^g x_2)$. For any $f_1, f_2 \in \mathrm{C}(X)$ we have

$$\lim_{g\to\mathsf{p}} \int f_1 \otimes f_2 \,\mathrm{d}(R^g\lambda) = \lim_{g\to\mathsf{p}} \int f_1 \cdot T^g f_2 \,\mathrm{d}\mu = \int \mathbb{E}(f_1|\mathsf{A}^\mu\mathbf{X}) \cdot \mathbb{E}(f_2|\mathsf{A}^\mu\mathbf{X}) \,\mathrm{d}\mu$$

by the above. Since the space of Baire probability measures on X^2 is a compact, Hausdorff space the sequence $g \mapsto R^g\lambda$ has a limit along p. Let ν be this limit. The above calculation implies that

$$\int f_1 \otimes f_2 \,\mathrm{d}\nu = \int \mathbb{E}(f_1|\mathsf{A}^\mu\mathbf{X}) \cdot \mathbb{E}(f_2|\mathsf{A}^\mu\mathbf{X}) \,\mathrm{d}\mu$$

for all $f_1, f_2 \in \mathrm{C}(X)$ as desired. □

Given a measure-preserving system $(\mathbf{X}, \mu)$ in $\mathcal{C}^{\mathrm{meas}}(G^2)$, the measure ν obtained by applying Lemma 6.1 to the measure-preserving system (X, T_2, μ) in $\mathcal{C}^{\mathrm{meas}}(G)$ is called the *relatively independent self-joining* of μ over $\mathsf{A}_2^\mu\mathbf{X}$. It follows immediately from (6.2) and the properties of conditional expectation that ν is invariant under the commuting G actions $R_1 = T_1 \times T_{12}$ and $R_2 = T_2 \times I$. Thus (X^2, R, ν) is a measure-preserving system in $\mathcal{C}^{\mathrm{meas}}(G^2)$. Lastly, writing π_1 and π_2 for the coordinate projections $X^2 \to X$, note that (6.2) implies that $\pi_1\nu = \pi_2\nu = \mu$ because all three measures agree on $\mathrm{C}(X)$.

We now turn to the proof of Theorem 1.8, which begins with the following lemma.

Lemma 6.3. *Let G be a discrete, minimally almost periodic group and p a minimal idempotent ultrafilter on G. Let $(\mathbf{X}, \mu)$ be a measure-preserving system in $\mathcal{C}^{\mathrm{meas}}(G^2)$ and let f_1 in $\mathrm{C}(X)$ be bounded by 1. Suppose that $(\mathbf{X}, \mu)$ is $(\varepsilon^2, f_1, \mathcal{I}_{1,2})$ sated. Then*

$$\left| \lim_{g\to\mathsf{p}} \int f_0 \cdot T_1^g f_1 \cdot T_{12}^g f_2 \,\mathrm{d}\mu - \lim_{g\to\mathsf{p}} \int f_0 \cdot T_1^g \mathbb{E}(f_1|\mathcal{I}_{1,2}\mathbf{X}) \cdot T_{12}^g f_2 \,\mathrm{d}\mu \right| < \varepsilon.$$

for any $f_0, f_2 \in \mathrm{C}(X)$ bounded by 1.

Proof. Define $u : G \to \mathrm{L}^2(\mathbf{X}, \mu)$ by $u(g) = T_1^g\phi \cdot T_{12}^g f_2$ where $\phi = f_1 - \mathbb{E}(f_1|\mathcal{I}_{1,2}\mathbf{X})$. We have

$$\lim_{g\to\mathsf{p}} \langle u(hg), u(g)\rangle = \lim_{g\to\mathsf{p}} \int (\phi \cdot T_1^h\phi) \cdot T_2^g(f_2 \cdot T_{12}^h f_2) \,\mathrm{d}\mu = \int \mathbb{E}(\phi \cdot T_1^h\phi|\mathsf{A}_2^\mu\mathbf{X}) \cdot \mathbb{E}(f_2 \cdot T_{12}^h f_2|\mathsf{A}_2^\mu\mathbf{X}) \,\mathrm{d}\mu$$

for every $h \in G$ by Theorem 5.1.

Let ν be the relatively independent self-joining of μ over $\mathsf{A}_2^\mu\mathbf{X}$, which exists by Lemma 6.1. Write π_1 and π_2 for the coordinate projections $X \times X \to X$. Define $\mathbf{Y} = (X \times X, R)$ where $R_1 = T_1 \times T_{12}$ and $R_2 = T_2 \times I$. We know from the above that $(\mathbf{Y}, \nu)$ is a $G \times G$ system and that $\pi_1 : (\mathbf{Y}, \nu) \to (\mathbf{X}, \mu)$ is a factor map. Another application of Theorem 5.1 yields

$$\lim_{h\to\mathsf{p}} \lim_{g\to\mathsf{p}} \langle u(hg), u(g)\rangle = \lim_{h\to\mathsf{p}} \int (\phi \otimes f_2) \cdot R_1^h(\phi \otimes f_2) \,\mathrm{d}\nu = \int \phi \circ \pi_1 \cdot f_2 \circ \pi_2 \cdot \mathbb{E}(\phi \otimes f_2|\mathsf{A}_1^\nu\mathbf{Y}) \,\mathrm{d}\nu. \tag{6.4}$$

We have $f_2 \circ \pi_2 \in \mathrm{C}(\mathcal{I}_2\mathbf{Y})$. Passing to a further extension $\mathbf{Z} \to \mathbf{Y}$ we may assume that $\mathbb{E}(\phi \otimes f_2|\mathsf{A}_1^\nu\mathbf{Z})$ has a representative in $\mathrm{C}(\mathcal{I}_1\mathbf{Z})$. By definition of $(\varepsilon^2, f_1, \mathcal{I}_{1,2})$ satedness it follows that (6.4) is bounded by

$$\varepsilon^2\|f_2\|_\infty\|\phi \otimes f_2\|_2 \leq \varepsilon^2\|f_2\|_\infty\|\phi\|_2\|f_2\|_\infty \leq \varepsilon^2,$$

so

$$\left| \lim_{g\to\mathsf{p}} \int f_0 \cdot T_1^g(f_1 - \mathbb{E}(f_1|\mathcal{I}_{1,2}\mathbf{X})) \cdot T_{12}^g f_2 \,\mathrm{d}\mu \right| \leq \varepsilon$$

by Proposition 4.4. □

Theorem 6.5. *Let G be a minimally almost periodic group and let $(\mathbf{X}, \mu)$ be a measure-preserving system in $\mathcal{C}^{\mathrm{erg}}(G^2)$. For any $\varepsilon > 0$ and any f_1 in $\mathrm{L}^\infty(X, \mu)$ bounded by 1 there is an extension $\pi : (\mathbf{Y}, \nu) \to (\mathbf{X}, \mu)$ in $\mathcal{C}^{\mathrm{erg}}(G^2)$ such that*

$$\left| \lim_{g\to\mathsf{p}} \int f_0 \cdot T_1^g f_1 \cdot T_{12}^g f_2 \,\mathrm{d}\mu - \lim_{g\to\mathsf{p}} \int \mathbb{E}(f_0 \circ \pi|\mathsf{A}_{1,12}^\nu\mathbf{Y}) \cdot T_1^g \mathbb{E}(f_1 \circ \pi|\mathcal{I}_{1,2}\mathbf{Y}) \cdot T_{12}^g \mathbb{E}(f_2 \circ \pi|\mathsf{A}_{12,2}^\nu\mathbf{Y}) \,\mathrm{d}\nu \right| < \varepsilon$$

for all minimal idempotent ultrafilters p on G and all f_0, f_2 in $\mathrm{L}^\infty(X, \mu)$ bounded by 1.

Proof. Let $(\mathbf{X},\mu)$ be a measure-preserving system in $\mathcal{C}^{\mathrm{erg}}(G^2)$ and fix f_1 in $\mathrm{L}^\infty(\mathbf{X},\mu)$ bounded by 1. Fix $\varepsilon>0$. By Theorem 3.6 and Proposition 3.10 we can find an extension $(\mathbf{Y},\nu)$ of $(\mathbf{X},\mu)$ in $\mathcal{C}^{\mathrm{erg}}(G^2)$ via a factor map π that is $(\varepsilon^2,f_1,\mathcal{I}_{1,2})$ sated in $\mathcal{C}^{\mathrm{meas}}(G^2)$. Lemma 6.3 implies that

$$\left|\lim_{g\to\mathsf{p}}\int f_0\cdot T_1^g f_1\cdot T_{12}^g f_2\,\mathrm{d}\mu-\lim_{g\to\mathsf{p}}\int (f_0\circ\pi)\cdot T_1^g\mathbb{E}(f_1\circ\pi|\mathcal{I}_{1,2}\mathbf{Y})\cdot T_{12}^g(f_2\circ\pi)\,\mathrm{d}\nu\right|<\varepsilon \tag{6.6}$$

for any f_0,f_2 in $\mathrm{C}(X)$ bounded by 1. By density and linearity we may assume $\mathbb{E}(f_1\circ\pi|\mathcal{I}_{1,2}\mathbf{Y})=h_1h_2$ where h_i is $\mathsf{A}_i^\nu\mathbf{Y}$-measurable. Under this assumption the second term in (6.6) becomes

$$\lim_{g\to\mathsf{p}}\int (f_0\circ\pi)h_1\cdot T_{12}^g((f_2\circ\pi)h_2)\,\mathrm{d}\nu=\int (f_0\circ\pi)h_1\cdot\mathbb{E}((f_2\circ\pi)h_2|\mathsf{A}_{12}^\nu\mathbf{Y})\,\mathrm{d}\nu$$

by Theorem 5.1. It follows that the above limit vanishes if $f_0\circ\pi\perp\mathsf{A}_{1,12}^\nu\mathbf{Y}$. Since conditional expectation is self-adjoint, the above limit also vanishes if $f_2\circ\pi\perp\mathsf{A}_{2,12}^\nu\mathbf{Y}$, which gives the desired result. □

7. Strong Recurrence

In this section we prove Theorem 1.9, which follows from Theorem 7.1 below by passing to a continuous model as follows. Theorem 7.1 is a version of [BCRZK16, Corollary 4.9] for limits along minimal idempotent ultrafilters. Given an action S of a group G on a probability space $(X,\mathscr{B},\mu)$ by measurable, measure-preserving maps, consider the space $\mathcal{A}$ of all bounded, measurable functions on $(X,\mathscr{B})$. Equipped with the supremum norm it becomes a C*-algebra. Let Ω be the Gelfand spectrum of $\mathcal{A}$. Then $\mathrm{C}(\Omega)$ and $\mathcal{A}$ are isomorphic as C*-algebras. The action S of G on $(X,\mathscr{B},\mu)$ induces an action T of G on Ω by continuous maps. Moreover, the measure μ induces a bounded, linear functional on $\mathrm{C}(\Omega)$ that can be identified with a Baire probability measure on Ω that is T-invariant.

Theorem 7.1. *Let G be a minimally almost periodic group and let $(\mathbf{X},\mu)$ be a measure-preserving system in $\mathcal{C}^{\mathrm{erg}}(G^2)$. For any minimal idempotent ultrafilter p in βG we have*

$$\lim_{g\to\mathsf{p}}\int f_0\cdot T_1^g f_1\cdot T_{12}^g f_2\,\mathrm{d}\mu\geq\left(\int f_0^{1/4}f_1^{1/4}f_2^{1/4}\,\mathrm{d}\mu\right)^4$$

for any non-negative measurable functions f_0,f_1,f_2 on X.

Proof. By the monotone convergence theorem we may assume that the functions f_0,f_1,f_2 are bounded. By Theorem 6.5 it suffices to show that

$$\lim_{g\to\mathsf{p}}\int \mathbb{E}(f|\mathsf{A}_{1,12}^\mu\mathbf{X})\cdot T_1^g\mathbb{E}(f|\mathcal{I}_{1,2}\mathbf{X})\cdot T_{12}^g\mathbb{E}(f|\mathsf{A}_{12,2}^\mu\mathbf{X})\,\mathrm{d}\mu\geq\left(\int f_0^{1/4}f_1^{1/4}f_2^{1/4}\,\mathrm{d}\mu\right)^4$$

for any ergodic measure-preserving system $(\mathbf{X},\mu)$ in $\mathcal{C}^{\mathrm{erg}}(G^2)$ and any non-negative $f_0,f_1,f_2\in\mathrm{L}^\infty(X,\mu)$. We prove that

$$g\mapsto\int \mathbb{E}(f_0|\mathsf{A}_{1,12}^\mu\mathbf{X})\cdot T_1^g\mathbb{E}(f_1|\mathcal{I}_{1,2}\mathbf{X})\cdot T_{12}^g\mathbb{E}(f_2|\mathsf{A}_{12,2}^\mu\mathbf{X})\,\mathrm{d}\mu \tag{7.2}$$

does not depend on g. Since any $\mathsf{A}_{1,2}^\mu\mathbf{X}$ measurable function can be approximated by linear combinations of functions of the form $\xi_1\xi_2$ where ξ_i is $\mathsf{A}_i^\mu\mathbf{X}$ measurable, we may replace f_0 above with $h_{0,1}h_{0,12}$, where $h_{j,i}$ is $\mathsf{A}_i^\mu\mathbf{X}$ measurable, and similarly f_1 and f_2 by $h_{1,1}h_{1,2}$ and $h_{2,12}h_{2,2}$ respectively. Then the above integral equals

$$\int h_{0,12}h_{2,12}\cdot h_{0,1}h_{1,1}\cdot T_{12}^g(h_{1,2}h_{2,2})\,\mathrm{d}\mu$$

and it therefore suffices to show that the sub-σ-algebras $\mathsf{A}_1^\mu\mathbf{X}$, $\mathsf{A}_2^\mu\mathbf{X}$, and $\mathsf{A}_{12}^\mu\mathbf{X}$ are jointly independent. Let now $h_i\in\mathrm{L}^\infty(X,\mu)$ be $\mathsf{A}_i^\mu\mathbf{X}$ measurable for $i=1,2,12$. We have to show

$$\int h_1\cdot h_2\cdot h_{12}\,\mathrm{d}\mu=\int h_1\,\mathrm{d}\mu\int h_2\,\mathrm{d}\mu\int h_{12}\,\mathrm{d}\mu.$$

From

$$\int h_1\cdot h_2\cdot h_{12}\,\mathrm{d}\mu=\int \mathbb{E}(h_1|\mathsf{A}_1^\mu\mathbf{X})\cdot\mathbb{E}(h_2|\mathsf{A}_2^\mu\mathbf{X})\cdot h_{12}\,\mathrm{d}\mu=\int h_1\cdot h_2\cdot\mathbb{E}(h_{12}|\mathsf{A}_{1,2}^\mu\mathbf{X})\,\mathrm{d}\mu$$

and the fact that the conditional expectation onto $\mathsf{A}^\mu_{1,2}\mathbf{X}$ commutes with T_{12} (since this σ-algebra is T_{12}-invariant), we may additionally assume that h_{12} is $\mathsf{A}^\mu_{1,2}\mathbf{X}$ measurable.

Ergodicity of $(\mathbf{X},\mu)$ implies the sub-σ-algebras $\mathsf{A}^\mu_1\mathbf{X}$ and $\mathsf{A}^\mu_2\mathbf{X}$ are independent. Thus

$$\mathrm{L}^2(\mathsf{A}^\mu_{1,2}\mathbf{X}) \cong \mathrm{L}^2(\mathsf{A}^\mu_1\mathbf{X}) \otimes \mathrm{L}^2(\mathsf{A}^\mu_2\mathbf{X})$$

and, thinking of h_{12} as an element of the right-hand side, we see that h_{12} is a member of

$$\mathrm{L}^2(\mathsf{C}^\mu_2\mathsf{A}^\mu_1\mathbf{X}) \otimes \mathrm{L}^2(\mathsf{C}^\mu_1\mathsf{A}^\mu_2\mathbf{X})$$

by Lemma 2.5. (Here C^μ_i denotes the σ-algebra of T_i almost periodic functions.) Since G is minimally almost periodic $\mathsf{C}_2\mathsf{A}^\mu_1\mathbf{X}$ and $\mathsf{C}_1\mathsf{A}^\mu_2\mathbf{X}$ are both the trivial σ-algebra. This proves (7.2) is constant. Finally

$$\begin{aligned}
&\lim_{g\to\mathsf{p}} \int \mathbb{E}(f_0|\mathsf{A}^\mu_{1,12}\mathbf{X}) \cdot T_1^g \mathbb{E}(f_1|\mathcal{I}_{1,2}\mathbf{X}) \cdot T_{12}^g \mathbb{E}(f_2|\mathsf{A}^\mu_{12,2}\mathbf{X})\,\mathrm{d}\mu \\
&= \int \mathbb{E}(f_0|\mathsf{A}^\mu_{1,12}\mathbf{X}) \cdot \mathbb{E}(f_1|\mathcal{I}_{1,2}\mathbf{X}) \cdot \mathbb{E}(f_2|\mathsf{A}^\mu_{12,2}\mathbf{X})\,\mathrm{d}\mu \\
&\geq \left(\int f_0^{1/4} f_1^{1/4} f_2^{1/4}\,\mathrm{d}\mu\right)^4
\end{aligned}$$

by Lemma A.1. □

8. Combinatorial results

We begin by proving Theorem 1.12. Fix an ultrafilter on $\mathbb{N}$ for taking ultraproducts. Let $n \mapsto G_n$ be a quasirandom sequence of finite groups. Let G be the ultraproduct of the sequence $n \mapsto G_n$ and let Ω be the ultraproduct of the sequence $n \mapsto G_n \times G_n$. We consider the commuting actions L_1 and L_2 of G on Ω defined by $L_1^g(x,y) = (gx,y)$ and $L_2^g(x,y) = (x,gy)$ respectively. The induced $G\times G$ action L is just the action of Ω on itself by left multiplication. We first note that this action is ergodic with respect to the Loeb measure m on Ω provided G is minimally almost periodic.

Lemma 8.1. *If G is minimally almost periodic then the $G\times G$ action L on (Ω,m) is ergodic.*

Proof. Given a Loeb measurable subset B of Ω we have $\mathrm{m}(B\cap(L^x)^{-1}B) = \mathrm{m}(B)^2$ for Loeb almost every $x\in\Omega$ by [BT14, Lemma 33]. Thus if B is almost invariant then $\mathrm{m}(B)\in\{0,1\}$. □

In fact [BT14, Lemma 33] implies L is weak mixing, but we will not need this.

It will be convenient later to pass to a model of this action on a compact, Hausdorff space. We do so by considering the C* algebra of bounded, measurable functions on Ω, which can be represented as $\mathrm{C}(X)$ for some compact, Hausdorff topological space X. Let $\mathscr{B}$ be the Baire sub-σ-algebra of X. The Loeb measure m on Ω passes to a probability measure μ on $(X,\mathscr{B})$ and the G actions L_1 and L_2 become actions T_1 and T_2 of G on X by homeomorphisms. Write T for the induced $G\times G$ action $(g_1,g_2)\mapsto T^{(g_1,g_2)}$ on X. Thus we have a measure-preserving system $(\mathbf{X},\mu)$ in $\mathcal{C}^{\mathrm{meas}}(G^2)$. This system is ergodic when G is minimally almost periodic.

Proposition 8.2. *If G is minimally almost periodic then $(\mathbf{X},\mu)$ is ergodic.*

Proof. Fix a continuous function f on X and a minimal idempotent ultrafilter p on $G\times G$. By Theorem 5.1 we have

$$\lim_{g\to\mathsf{p}} \int \phi\cdot T^g f\,\mathrm{d}\mu = \int \phi\cdot\mathbb{E}(f|\mathsf{A}^\mu\mathbf{X})\,\mathrm{d}\mu$$

for any continuous ϕ. By evaluating the left hand side on Ω rather than on X, we obtain

$$\lim_{g\to\mathsf{p}} \int \phi\cdot T^g f\,\mathrm{d}\mu = \int \phi\,\mathrm{d}\mu \int f\,\mathrm{d}\mu$$

by Theorem 5.1 and the previous lemma. Thus $\mathbb{E}(f|\mathsf{A}^\mu\mathbf{X}) = \int f\,\mathrm{d}\mu$. The same is true for any function that can be approximated in $\mathrm{L}^2(X,\mathscr{B},\mu)$ by continuous functions. □

Now we turn to the proof of Theorem 1.12.

Proof of Theorem 1.12. Suppose that the conclusion fails for some $0 < \alpha < 1$ and some $\varepsilon > 0$. Then we can find sequences $D_n \to \infty$ and $K_n \to \infty$ in $\mathbb{N}$, a sequence $n \mapsto G_n$ of D_n-quasirandom groups in $\mathcal{F}$, and sets $A_n \subset G_n \times G_n$ with $|A_n| \geq \alpha |G_n|^2$ such that

$$R_n := \left\{ g \in G_n : \frac{|A_n \cap (1,g)^{-1} A_n \cap (g,g)^{-1} A_n|}{|G_n|^2} > \alpha^4 - \varepsilon \right\}$$

is not right K_n-syndetic. Since being right K-syndetic for some finite K is a first order property, it follows that the ultraproduct R of the sequence $n \mapsto R_n$ is not right K-syndetic for any K, and therefore not right syndetic.

Let G be the ultraproduct of the sequence $n \mapsto G_n$ and let Ω be the ultraproduct of the sequence $n \mapsto G_n \times G_n$. Since $\mathcal{F}$ is a quasirandom ultraproduct class the group G is minimally almost periodic. Let A be the internal subset of Ω determined by the sequence $n \mapsto A_n$. We have

$$R \supseteq \{ g \in G : \mu(A \cap (1,g)^{-1} A \cap (g,g)^{-1} A) > \alpha^4 - \varepsilon/2 \}$$

where μ is the Loeb measure on Ω. But R is right central* by Theorem 7.1, and therefore right syndetic by Lemma 4.6, giving the desired contradiction. □

Proof of Theorem 1.15. Fix a quasirandom sequence $n \mapsto G_n$ with $G_n \hookrightarrow G_{n+1}$ for all $n \in \mathbb{N}$. Suppose the theorem is false for some $0 < \alpha < 1$ and some $\varepsilon > 0$. Then we have a sequence of sets $A_n \subset G_n \times G_n$ with $|A_n| \geq \alpha |G_n|^2$ and

$$|\{ g \in G_n : A_n \cap (1,g)^{-1} A_n \cap (g,g)^{-1} A_n \neq \varnothing \}| \leq (1-\varepsilon)|G_n| \tag{8.3}$$

for all $n \in \mathbb{N}$.

Let G be the direct limit of the sequence $n \mapsto G_n$. Put $A = \cup \{A_n : n \in \mathbb{N}\}$ in $G \times G$. We have

$$\limsup_{N \to \infty} \frac{|A \cap G_N \times G_N|}{|G_N \times G_N|} \geq \alpha$$

so A has positive upper density with respect to the Følner sequence $N \mapsto G_N \times G_N$ in $G \times G$. By [CZK15, Lemma 5.2] there is an ergodic action T of $G \times G$ on a compact, metric probability space $(X, \mathscr{B}, \mu)$, an open set $U \subset X$ with $\mu(U) = \overline{\mathrm{d}}_\Phi(A)$, and a Følner sequence Ψ on $G \times G$ such that

$$\overline{\mathrm{d}}_\Psi((g_1,h_1)^{-1} A \cap \cdots \cap (g_n,h_n)^{-1} A) \geq \mu((T^{(g_1,h_1)})^{-1} U \cap \cdots \cap (T^{(g_n,h_n)})^{-1} U)$$

for all $g_n, h_n \in G$. In particular

$$\overline{\mathrm{d}}_\Psi(A \cap (1,g)^{-1} A \cap (g,g)^{-1} A) \geq \mu(U \cap (T_1^g)^{-1} U \cap (T_1^g T_2^g)^{-1} U)$$

for all $g \in G$. It follows from [BCRZK16, Corollary 4.9] that for every $\varepsilon > 0$ the set

$$\{ g \in G : \overline{\mathrm{d}}_\Psi(A \cap (1,g)^{-1} A \cap (g,g)^{-1} A) > 0 \}$$

has full density with respect to every Følner sequence in G. In particular, it has full density with respect to the Følner sequence $N \mapsto G_N \times G_N$, contradicting (8.3) for n large enough. □

Appendix A. Chu's inequality

We use the following slightly generalized version of Chu's lower bound for a product of conditional expectations [Chu11, Lemma 1.6].

Lemma A.1. *Let $f_0, \ldots, f_n$ be non-negative integrable functions on a probability space (X, B, μ) and let $B_1, \ldots, B_n \subset B$ be arbitrary sub-σ-algebras. Then*

$$\int f_0 \prod_{i=1}^n \mathbb{E}(f_i | B_i) \geq \Big(\int \prod_{i=0}^n f_i^{\frac{1}{n+1}} \Big)^{n+1}.$$

The main advantage of the present formulation is the ability to take $f_0 \equiv 1$.

Proof. Note that $\{f_i > 0\} \subset \{\mathbb{E}(f_i|B_i) > 0\}$ up to a set of measure zero for every i. By Hölder's inequality we have

$$\begin{aligned}\int \prod_{i=0}^{n} f_i^{\frac{1}{n+1}} &= \int \big(f_0^{\frac{1}{n+1}} \prod_{i=1}^{n} \mathbb{E}(f_i|B_i)^{\frac{1}{n+1}}\big) \cdot \prod_{i=1}^{n} \Big(1_{\{\mathbb{E}(f_i|B_i)>0\}} \frac{f_i}{\mathbb{E}(f_i|B_i)}\Big)^{\frac{1}{n+1}} \\ &\leq \Big(\int f_0 \prod_{i=1}^{n} \mathbb{E}(f_i|B_i)\Big)^{\frac{1}{n+1}} \cdot \prod_{i=1}^{n} \Big(\int_{\{\mathbb{E}(f_i|B_i)>0\}} \frac{f_i}{\mathbb{E}(f_i|B_i)}\Big)^{\frac{1}{n+1}}.\end{aligned}$$

Since the functions $\mathbb{E}(f_i|B_i)^{-1}$ are B_i-measurable, f_i's may be replaced by their expectations onto B_i in the integrals in the second factor. Thus we obtain

$$\begin{aligned}&= \Big(\int f_0 \prod_{i=1}^{n} \mathbb{E}(f_i|B_i)\Big)^{\frac{1}{n+1}} \cdot \prod_{i=1}^{n} |\{\mathbb{E}(f_i|B_i) > 0\}|^{\frac{1}{n+1}} \\ &\leq \Big(\int f_0 \prod_{i=1}^{n} \mathbb{E}(f_i|B_i)\Big)^{\frac{1}{n+1}}.\end{aligned}$$

□

It would be interesting to know whether the lower bound in Lemma A.1 is sharp for some characteristic functions f_i. There are two sources of inefficiency in its proof: the Hölder inequality and the estimate $|\{\mathbb{E}(f_i|B_i)\}| \leq 1$. It is clear that the second source of inefficiency can be easily eliminated. On the other hand, an example in which the Hölder inequality gives a sharp estimate can be found in [CZK15, Appendix B] (in hindsight this provides an explanation for why that example, which has been initially found numerically, works). However, it is not clear whether both sources of inefficiency can be controlled simultaneously.

References

[Aus10] T. Austin. "On the norm convergence of non-conventional ergodic averages". In: *Ergodic Theory Dynam. Systems* 30.2 (2010), pp. 321–338.

[Aus13] T. Austin. "Non-conventional ergodic averages for several commuting actions of an amenable group". In: *J. Analyse Math.* (2013). To appear. arXiv:1309.4315 [math.DS].

[Aus15] T. Austin. "Quantitative equidistribution for certain quadruples in quasi-random groups". In: *Combin. Probab. Comput.* 24.2 (2015). With erratum, pp. 376–381. arXiv:1310.6781 [math.CO].

[Aus16] T. Austin. "Ajtai–Szemerédi Theorems over quasirandom groups". In: *Recent Trends in Combinatorics*. 2016. arXiv:1503.08746 [math.CO].

[BCRZK16] V. Bergelson, C. Christopherson, D. Robertson, and P. Zorin-Kranich. "Finite products sets and minimally almost periodic groups". In: *J. Func. Anal* 270.6 (2016), pp. 2126–2167. arXiv:1402.4736 [math.DS].

[Ber00] V. Bergelson. "Ergodic theory and Diophantine problems". In: *Topics in symbolic dynamics and applications (Temuco, 1997)*. Vol. 279. London Math. Soc. Lecture Note Ser. Cambridge Univ. Press, Cambridge, 2000, pp. 167–205.

[Ber03] V. Bergelson. "Minimal idempotents and ergodic Ramsey theory". In: *Topics in dynamics and ergodic theory*. Vol. 310. London Math. Soc. Lecture Note Ser. Cambridge Univ. Press, Cambridge, 2003, pp. 8–39.

[BF09] V. Bergelson and H. Furstenberg. "WM groups and Ramsey theory". In: *Topology Appl.* 156.16 (2009), pp. 2572–2580.

[BH90] V. Bergelson and N. Hindman. "Nonmetrizable topological dynamics and Ramsey theory". In: *Trans. Amer. Math. Soc.* 320.1 (1990), pp. 293–320.

[BM07] V. Bergelson and R. McCutcheon. "Central sets and a non-commutative Roth theorem". In: *Amer. J. Math.* 129.5 (2007), pp. 1251–1275.

[BMZ97] V. Bergelson, R. McCutcheon, and Q. Zhang. "A Roth theorem for amenable groups". In: *Amer. J. Math.* 119.6 (1997), pp. 1173–1211.

[BS85] L. Babai and V. T. Sós. "Sidon sets in groups and induced subgraphs of Cayley graphs". In: *European J. Combin.* 6.2 (1985), pp. 101–114.

[BT14] V. Bergelson and T. Tao. "Multiple Recurrence in Quasirandom Groups". In: *Geom. Funct. Anal.* 24.1 (2014), pp. 1–48. arXiv:1211.6372 [math.DS].

[Chu11] Q. Chu. "Multiple recurrence for two commuting transformations". In: *Ergodic Theory Dynam. Systems* 31.3 (2011), pp. 771–792. arXiv:0912.3381 [math.DS].

[Cut83] N. J. Cutland. "Nonstandard measure theory and its applications". In: *Bull. London Math. Soc.* 15.6 (1983), pp. 529–589.

[CZK15] Q. Chu and P. Zorin-Kranich. "Lower bound in the Roth theorem for amenable groups". In: *Ergodic Theory Dynam. Systems* 35.6 (2015), pp. 1746–1766. arXiv:1309.6095 [math.DS].

[Ell58] R. Ellis. "Distal transformation groups". In: *Pacific J. Math.* 8 (1958), pp. 401–405.

[Fur77] H. Furstenberg. “Ergodic behavior of diagonal measures and a theorem of Szemerédi on arithmetic progressions”. In: *J. Analyse Math.* 31 (1977), pp. 204–256.

[Fur81] H. Furstenberg. “Recurrence in ergodic theory and combinatorial number theory”. M. B. Porter Lectures. Princeton University Press, Princeton, N.J., 1981, pp. xi+203.

[Gow08] W. T. Gowers. “Quasirandom groups”. In: *Combin. Probab. Comput.* 17.3 (2008), pp. 363–387. arXiv:0710.3877 [math.CO].

[LG61] K. de Leeuw and I. Glicksberg. “Applications of almost periodic compactifications”. In: *Acta Math.* 105 (1961), pp. 63–97.

[Phe01] R. R. Phelps. “Lectures on Choquet’s theorem”. Second. Vol. 1757. Lecture Notes in Mathematics. Springer-Verlag, Berlin, 2001, pp. viii+124.

[Rob96] D. J. S. Robinson. “A course in the theory of groups”. Second. Vol. 80. Graduate Texts in Mathematics. Springer-Verlag, New York, 1996, pp. xviii+499.

[Sch07] C. Schnell. “Idempotent ultrafilters and polynomial recurrence”. Preprint. 2007. arXiv:0711.0484 [math.DS].

[Tse14] A. Tserunyan. “Mixing and triple recurrence in probability groups”. Preprint. 2014. arXiv:1405.5629 [math.DS].

[Yan16] Y. Yang. “Ultraproducts of quasirandom groups with small cosocles”. In: *J. Group Theory* (2016). To appear. arXiv:1405.6276v2 [math.GR].